Persistent Bias Among Local Election Officials *

D. Alex Hughes[1], Micah Gell-Redman[2], Charles Crabtree[3], Natarajan Krishnaswami[1], Diana Rodenberger[1] and Guillermo Monge[1]

[1]School of Information, University of California, Berkeley
[2]Department of International Affairs and Department of Health Policy & Management, University of Georgia
[3]Department of Political Science, University of Michigan

April 11, 2019

Abstract

Results of an audit study conducted during the 2016 election cycle demonstrate that bias toward Latinos observed during the 2012 election has persisted. In addition to replicating previous results, we show that Arab/Muslim Americans face an even greater barrier to communicating with local election officials, but we find no evidence of bias toward blacks. An innovation of our design allows us to measure whether emails were opened by recipients, which we argue provides a direct test of implicit discrimination. We find evidence of implicit bias toward Arab/Muslim senders only.

*The data, code, and compute environment required to replicate all analyses in this article are available at the Journal of Experimental Political Science Dataverse within the Harvard Dataverse Network, at: https://doi.org/10.7910/DVN/8E1IIM (Hughes et al., 2019). The authors are aware of no conflicts of interest regarding this research.

Racial bias that limits access to the ballot threatens basic principles of democratic equality. One potential source of bias that has received little attention are the street level bureaucrats who administer elections in the U.S. (Lipsky, 1980). An audit study conducted during the 2012 U.S. election cycle showed these local election officials responded at significantly lower rates to inquiries from voters with putatively Latino, as opposed to white, surnames (White, Nathan and Faller, 2015). In this paper we report the results of a similar audit study performed during the 2016 election cycle. We find that the previously observed bias against Latinos is persistent. We also extend the previous study by testing the effects of two racial primes other than Latino. Voters with Arab/Muslim names received responses at significantly lower rates (11 percentage points) than whites, while black voters did not.

The two primary motivations for this study are to determine whether the previous finding of bias toward Latinos stands up to replication, and to examine whether this bias extends to blacks and Arab/Muslim Americans. In spite of the ample evidence of racial disparities in political participation (Hajnal and Lee, 2011; Abrajano and Alvarez, 2010; Hajnal and Abrajano, 2015; García-Bedolla and Michelson, 2012) and in every-day life (Bertrand and Mullainathan, 2004), relatively little empirical work demonstrates the role of race in limiting access to the ballot in contemporary America (McNulty, Dowling and Ariotti, 2009), and some claims in this area have aroused skepticism (Hajnal, Lajevardi and Nielson, 2017; Grimmer et al., 2018). The pervasive discrimination that blacks face in various arenas of American politics (Butler, 2014) suggests that this group could be at risk of bias in interacting with local election officials. While there is also ample evidence of discrimination toward Arab and Muslim Americans (Gaddis and Ghoshal, 2015), this group has received comparatively less attention from scholars (Jamal and Naber, 2007; Panagopoulos, 2006). In an era of political rhetoric increasingly characterized by appeals to group identity, it is particularly important to understand how racially-motivated bias impacts the day-to-day mechanics of elections for a range of racial/ethnic groups.

To seek evidence of bias, we focus on the thousands of local-level administrators charged with conducting elections in the United States. These bureaucrats are generally capable of exercising discretion in carrying out their job duties, which include responding to inquiries about the mechanics of voting and eligibility to participate in elections. Our core contention is that in exercising such discretion, street-level bureaucrats may be consciously or unconsciously influenced by the characteristics (e.g., race or partisanship) of individuals seeking public services (Lipsky, 1980; White, Nathan and Faller, 2015).

1 Experiment Design

To determine the extent to which previously documented bias is persistent and extends to other racial groups, we conduct an email audit study of local election officials (Pager, 2003).[1] Our intended sample comprises all such officials with publicly available email addresses and the analytic sample includes 6,439 local election officials from 44 states (Figure A1).

The experimental stimulus consists of a single email sent to each local election official. All emails follow the same structure, greeting the official by name, referencing voter identification laws, and asking about the requirements to vote in the state corresponding to the official. Our design closely parallels White, Nathan and Faller (2015), but differs in that we send only messages that mention voter ID laws. Additionally, to minimize possible spillover issues, we create 27 variants of this request (SI section A4 and section A6).

Our experimental treatment is the putative identity of the email sender. In line with convention we expose officials to four distinct group identities by manipulating senders' names (Bertrand and Mullainathan, 2004; Bertrand and Duflo, 2017; Butler and Homola, 2017). Because the identities signaled in our treatments have elements which could be described as racial, ethnic, or religious, we refer to these generically as group identity treatments. To mitigate possible name effects, each group identity condition is signaled by 100 unique names. We check that the chosen names reliably prime ethnicity by conducting a manipulation check on Amazons Mechanical Turk service in which workers read sets of names and ascribe probabilities that a name belongs to a particular racial or ethnic group.[2] In total, we send 4,900 unique experimental conditions which combine variants of the contact language with treatment identities.

1.1 Treatment assignment and implementation

We block treatment assignment on logged population density, two-party vote share in the 2012 presidential election, percent African American, percent Latino, percent of households with incomes below 150 percent of the federal poverty level, and a dummy variable indicating whether a county was previously covered by Section 5 of the Voting Rights Act. Further details are provided in SI section A8.

[1] We received Human Subjects approval from the University of California, Berkeley and University Michigan Human Subjects Committees. Both committees waived the requirement of informed consent. Additional implementation details are made available in the Supplemental Information. The study design, and pre-analysis plan were registered at Evidence in Governance and Politics. Data, code, and computing environments are available at ALEX INSERT.

[2] SI section A7 describes the procedure for choosing names, and section A17 provides the complete list of names.

Within each block we assign local election officials a racial condition and message version at random.

We sent 6,235 emails the morning of October 31, 2016, one email to each election official that was a part of the study.[3] Emails were sent from a purpose-built domain, `ez-webmail.com`. Sending addresses took the form of the senders' first initial, last name, and a two-digit string between twenty and forty. To mitigate the possibility that elections officials would be suspicious of our contact, we structured the email headers so that inboxes displayed the full name of the purported voter (see Figure A1). The variety in our treatments was intended to reduce the likelihood that different offices would receive emails from identical senders. In twenty-nine of the forty-three states in our analytic sample every official received a contact from a distinct name.

One key innovation in this experiment permits the identification of whether emails were received and opened by election officials. We include a 1x1 pixel image with a unique link – commonly referred to as a tracking pixel – in the email body so that upon opening the email, most email clients loaded the image from our server and provided a positive record that the email had been opened by a particular official. This measurement permits inference about differential open-rates, a test of implicit bias we examine in subsection 2.1.

An open question in correspondence studies concerns whether observed effects are merely an artifact of differential treatment of stimulus by the internet and email infrastructure, i.e., spam filters. Through pilot testing we are able to comment on this question. Before taking steps to develop positive server reputation, no messages reached any test inboxes. However, by carefully managing our digital authentication and consulting with individuals at a digital marketing company, in pilot testing we were able to place every message, from every attempted sender, into test inboxes (see SI section A2).

The choice to contact election officials eight days before the election is designed to make our study reflective of the real constraints on individuals seeking and providing information about voting requirements. To minimize the impact of our intervention on election officials' time, the specific request contained in the email is one that would require little effort to fulfill. Using data gathered via our mailing system, we estimate that the median time to compose and send a response to our email is three minutes, six seconds. We contend that any costs borne by public officials as a result of our intervention are counterbalanced by the benefits of uncovering persistent bias in electronic communications between constituents and local election officials.

[3] We also sent two waves of pilot email, 54 on October 26, 2016; and, 146 on October 28, 2016. For details, see SI section A12.

Table 1: Response Rates by Experimental Condition

Ethnic Cue	**White**	**Minority**	**Latino**	**Black**	**Arab**
Response Rate (%)	61.3	56.6	58.4	61.4	50.1
Standard Error	1.21	0.71	1.23	1.21	1.25
N	1,611	4,828	1,609	1,613	1,606

Notes: The *Minority* column includes all data from the *Latino*, *Black*, and *Arab* columns. Response rates and standard errors are reported in percentage terms.

Our pre-registered analysis uses a single outcome measure, GotResponse, coded 1 if an election official replied to our email prior to election day, and 0 otherwise.[4] We do not count auto-replies, away messages, or bounces as valid replies. We further report an exploratory analyses of a novel outcome measure made possible through our engineering: whether a local election official opened the message.

2 Results

Overall, 57.8 percent of the emails we sent received at least one reply from local election officials. While lower than the 67.7 percent response rate previously obtained from a similar sample (White, Nathan and Faller, 2015), this rate compares favorably with experiments on elected officials in the U.S., suggesting that our requests were taken at face value (Butler and Broockman, 2011).

Election officials respond at considerably lower rates when queries come from minority as opposed to white senders (difference in mean, $\Delta\mu = -4.70$ percentage points, *Wilcox Rank-Sum* $P < 2 \times 10^{-16}$). However, as we report in Table 1 responsiveness to minority senders is not uniformly lower. Nonparametric tests using white senders as the baseline find that a Latino name is sufficient to suppress the likelihood of a response by nearly 3 percentage points ($\Delta\mu = -2.97$, $P = 0.07$). Strikingly, an Arab/Muslim name lowers the likelihood of a response by greater than 11 percentage points ($\Delta\mu = -11.3$, $P < 1 \times 10^{-10}$). In contrast, black senders receive responses at a rate indistinguishable from white senders ($\Delta\mu = 0.11$, $P = 0.90$). Figure 1 (a) plots the Intent to Treat (ITT) causal effects of our treatments. Regression estimates with robust standard errors are reported in columns 1 and 2 of Table A6, and produce similar results.

Figure 1 (b) plots a precision weighted meta-analysis estimate (Gerber and

[4] Pre-analysis Plan Registered at EGAP (Hughes, Gell-Redman and Crabtree, 2016).

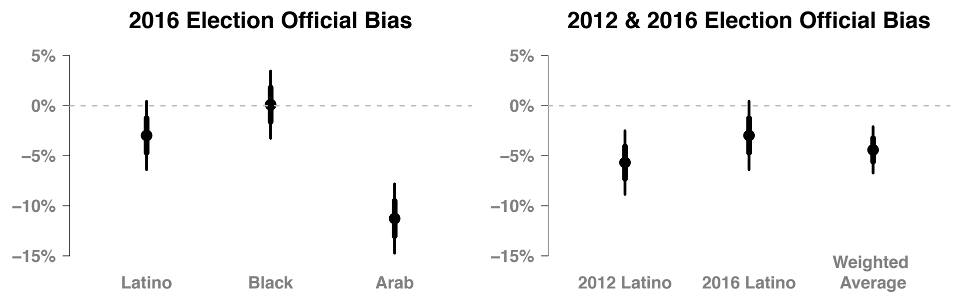

Figure 1: Points represent the ITT, the estimated difference in response rates to emails from the named identity, compared to the white response rate baseline. Thick bars report $ITT \pm SE$, thin bars report $ITT \pm 1.96 \times SE$. All estimates are difference in means, except the *Weighted Average* which estimates a precision weighted difference (Gerber and Green, 2012) utilizing 2012 (White, Nathan and Faller, 2015) and 2016 Latino evidence.

Green, 2012, p. 361) that combines the results of our intervention with those previously reported (White, Nathan and Faller, 2015). These data, gathered in independent audits conducted over two election cycles, show that Latinos receive replies from local election officials at a rate 4.4 percentage points lower than whites ($\Delta\mu = 4.4$, precision weighted $SE = 1.18$, $P < 0.0001$).

While the persistence of the treatment of Latino senders in the 2012 and 2016 elections is remarkable, perhaps more striking is the finding that Arab/Muslim names suffer a penalty more than two times greater than the one produced by a Latino stimulus. One potential concern is that the observed effect could be driven by the implausibility of the treatment, since many parts of the country do not have any appreciable population of Arab-Americans. To examine this possibility, we investigate whether treatment effects are smaller in the jurisdictions where Arab-Americans are more numerous. If treatment effects are driven by implausibility then they should be smaller in places where the presence of citizens with Arab names are more plausible. We do not find clear evidence that the proportion of Arab Americans moderates the treatment effect (Table A13, Model 3; Table A14; Table A15). Our most credible estimates find a 10.6 percentage point bias against Arab senders in counties with no Arab population ($\Delta\mu = -10.6$ percentage points, $SE = 2.5, P < 0.001$), but only a 2.6 percentage point improvement in the highest Arab population quartile of counties ($\delta\Delta\mu = +2.6$ percentage points, $SE = 4.4, P = 0.55$), although the distribution of Arab American settlement limits the strength of this robustness check.[5]

[5] In the highest Arab quartile, the mean Arab population is 1%.

2.1 Evidence of implicit discrimination

Local election officials who receive our intervention demonstrate bias insofar as they respond differentially based only on the signal of group identity delivered through our treatments. This observed response behavior is part of a chain of actions: the official must open, read, and then respond to the email. Standard analyses of audit experiments, which report an indicator of response or non-response as the dependent variable, focus only on the final result of this compound process. Innovations of our design allow us to consider the outcome at a prior step, the decision by the official to open the received email, conditional on the treatment delivered.

To respond to our experimental stimulus, an election official must identify our request from among the large number of other requests, categorize it mentally, and then open it. We argue that opening an email is a high-volume, low-attention task of the type scholars have associated with implicit, rather than explicit bias (Devine, 1989; Bertrand, Chugh and Mullainathan, 2005, p.96). The pattern of email opens suggests that, indeed, elections officials may be unintentionally or automatically screening requests from Arab/Muslim senders. There is no difference in open rates between white and latino names ($\Delta\mu = -0.74$ percentage points, $SE = 1.7$, $P = 0.68$) or white and black names ($\Delta\mu = -0.24$, $SE = 1.7$, $P = 0.90$). However, there is a pronounced gap in open rates for emails sent by senders with Arab/Muslim names, who have their emails opened at a rate 6.8 percentage points lower than white senders ($\Delta\mu = -6.8$, $SE = 1.8$, $P = 0.00013$).

2.2 Awareness of experiment

During the analysis phase of this project, it came to the researchers' attention another entity was pursuing a similar line of research using the same sending domain as White, Nathan and Faller (2015). As a result, some public officials became concerned that an audit study might be underway. News reports claim that these concerns prompted the National Association of Secretaries of State (NASS) to alert its state branches, who in turn had the opportunity to alert individual officials. In sum, some of our experimental subjects may have become aware of the presence of interventions.

Subjects' awareness of the intervention poses a general threat to audit studies, either by compromising independence between units, or by violating the exclusion restriction if minority names are more likely to raise suspicion than white names. Because subjects' awareness might prevent identification of causal effects, researchers should mitigate this risk by using many identities and a well-tuned sending architecture whenever feasible. When there is any observable information about the possibility of discovery, researchers can use this information to evaluate

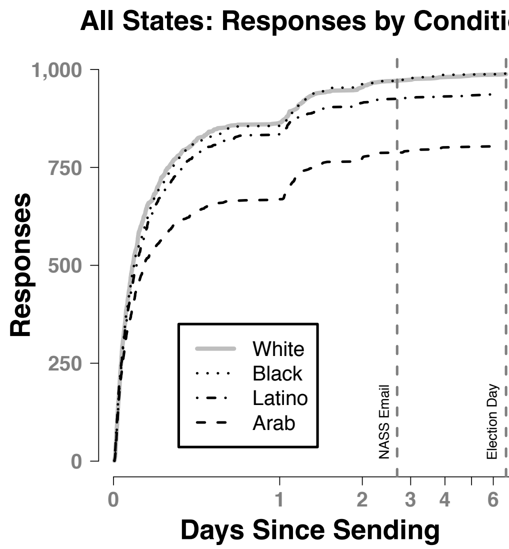

Figure 2: Rapidly slowing rates of response. The vertical axis plots the cumulative number of respones, split by group identity of sender; the horizontal axis plots time since sending. Election Day and NASS emails are noted with vertical dashed lines. Responses follow a clear diurnal rhythm, and patterns of bias appear rapidly.

whether apparent differences are likely the result of discovery.

Analysis of the timing of responses in this experiment does not suggest that discovery is leading to the observed results. First, as we present in Figure 2, the systematic pattern of unresponsiveness to minority names appears rapidly and well before the reported NASS broadcast. Second, as we report in Table A11 and Table A12, models that censor response data at the time of the NASS broadcast, and models that exclude states that witnessed interference between units both produce estimates very similar to our main results.

3 Conclusion

Previous experimental evidence showed local election officials were less responsive to inquiries from Latinos, raising concerns about bias in the electoral process. Using a similar experimental design, we demonstrate the firm basis for these concerns by replicating the initial finding. We also extend the results by testing for bias against other groups.

Our intervention showed Arab/Muslim Americans to be markedly disadvantaged in their interactions with local election officials. This finding is particularly salient given that it is not simply an artifact of Arab/Muslims being a relatively less numerous part of the electorate. We encountered no evidence of bias from local election officials toward African Americans, making ours at least the third recent study to produce a similarly unexpected null finding (Einstein and Glick, 2017; Gell-Redman et al., 2018). Rather than evidence of a lack of bias against African Americans, these null findings may be an artifact of the correspondence study method in which name alone, rather than other cues such as appearance, is used to signal identity.

Through this design, we also engage a challenge inherent to all audit studies, the risk that subjects become aware of the experiment. The relatively low technical sophistication required to conduct some forms of audit studies, mated with the potentially large sample size that is possible through email-based audits make these designs a potentially attractive way to identify discriminatory behavior. However, in an increasingly crowded field, researchers must face the possibility that experimental subjects become aware of the study, thereby damaging the inference. We determined that sending 4,900 distinct treatments on a custom-built server provided the best balance of a low possibility of discovery with the ability to identify a novel open rate outcome measure, and we would encourage future researchers to make a similar assessment.

References

Abrajano, Marisa A. and Michael M. Alvarez. 2010. *New Faces, New Voices: The Hispanic Electorate in America*. Princeton University Press.

Bertrand, M. and E. Duflo. 2017. Field Experiments on Discrimination. In *Handbook of Field Experiments*, ed. Abhijit Vinayak Banerjee and Esther Duflo. Vol. 1 North-Holland pp. 309 – 393.
URL: *http://www.sciencedirect.com/science/article/pii/S2214658X1630006X*

Bertrand, Marianne, Dolly Chugh and Sendhil Mullainathan. 2005. "Implicit discrimination." *American Economic Review* pp. 94–98.

Bertrand, Marianne and Sendhil Mullainathan. 2004. "Are Emily and Greg More Employable Than Lakisha and Jamal? A Field Experiment on Labor Market Discrimination." *American Economic Review* 94(4):991–1013.

Butler, Daniel M. 2014. *Representing the Advantaged: How Politicians Reinforce Inequality*. Cambridge University Press.

Butler, Daniel M and David E Broockman. 2011. "Do Politicians Racially Discriminate Against Constituents? A Field Experiment on State Legislators." *American Journal of Political Science* 55(3):463–477.

Butler, Daniel M and Jonathan Homola. 2017. "An Empirical Justification for the Use of Racially Distinctive Names to Signal Race in Experiments." *Political Analysis* 25(1):122–130.

Devine, Patricia G. 1989. "Stereotypes and prejudice: Their automatic and controlled components." *Journal of personality and social psychology* 56(1):5.

Einstein, Katherine Levine and David M. Glick. 2017. "Does Race Affect Access to Government Services? An Experiment Exploring Street-Level Bureaucrats and Access to Public Housing." *American Journal of Political Science* 61:100–116.

Gaddis, S Michael and Raj Ghoshal. 2015. "Arab American Housing Discrimination, Ethnic Competition, and the Contact Hypothesis." *The Annals of the American Academy of Political and Social Science* 660(1):282–299.

García-Bedolla, Lisa and Melissa R. Michelson. 2012. *Mobilizing Inclusion: Transforming the Electorate Through Get-out-the-Vote Campaigns*. Yale University Press.

Gell-Redman, Micah, Neil Visalvanich, Charles Crabtree and Christopher Fariss. 2018. "It's all about race: How state legislators respond to immigrant constituents." *Political Research Quarterly* .

Gerber, Alan S and Donald P Green. 2012. *Field experiments: Design, analysis, and interpretation*. WW Norton.

Grimmer, Justin, Eitan Hersh, Marc Meredith, Jonathan Mummolo and Clayton Nall. 2018. "Obstacles to estimating voter ID laws' effect on turnout." *Journal of Politics* 80(3).

Hajnal, Zoltan and Marisa Abrajano. 2015. *White Backlash: Immigration, Race, and American Politics*. Princeton University Press.

Hajnal, Zoltan, Nazita Lajevardi and Lindsay Nielson. 2017. "Voter Identification Laws and the Suppression of Minority Votes." *The Journal of Politics* 79(2):363–379.

Hajnal, Zoltan and T. Lee. 2011. *Why Americans Don't Join the Party: Race, Immigration, and the Failure (of Political Parties) to Engage the Electorate*. Princeton University Press.

Hughes, D. Alex, Micah Gell-Redman and Charles Crabtree. 2016. "Who gets to vote?" Evidence in Government and Politics, EGAP ID: 20161001AA.
URL: *http://egap.org/registration/2183*

Hughes, D. Alex, Micah Gell-Redman, Charles Crabtree, Natarajan Krishnaswami, Guillermo Monge and Diana Rodenberger. 2019. "Replication Data for: Persistent Bias Among Local Election Officials." *Harvard Dataverse* doi:10.7910/DVN/8E1IIM.

Jamal, Amaney and Nadine Naber. 2007. *Race and Arab Americans Before and After 9/11: From Invisible Citizens to Visible Subjects*. Syracuse University Press.

Lipsky, Michael. 1980. *Street-level Bureaucracy: Dilemmas of the Individual in Public Services*. Russell Sage.

McNulty, John E., Conor M. Dowling and Margaret H. Ariotti. 2009. "Driving Saints to Sin: How Increasing the Difficulty of Voting Dissuades Even the Most Motivated Voters." *Political Analysis* 17(4):435–455.

Pager, Devah. 2003. "The mark of a criminal record." *American journal of sociology* 108(5):937–975.

Panagopoulos, Costas. 2006. "The Polls-Trends: Arab and Muslim Americans and Islam in the Aftermath of 9/11." *Public Opinion Quarterly* 70(4):608–624.

White, Ariel R., Noah L. Nathan and Julie K. Faller. 2015. "What Do I Need to Vote? Bureaucratic Discretion and Discrimination by Local Election Officials." *American Political Science Review* 109(1):129–142.

Appendix for Persistant Bias Among Local Election Officials

March 5, 2019

Contents

A1 Email Scraping

We collected email and personal contact information from local election officials by programmatically visiting state-maintained sites of local election official contact information. We do not include the following states' local election officials in our assignment to treatment: Alaska, Hawaii, Maine, Maryland, Missouri, and New Jersey. We exclude Alaska because local election official jurisdictions were not mappable onto census area delineations for covariate data. We exclude Hawaii because a single board member represented each island, and the state did not provide individual email addresses for each island; rather, there was a single catch-all address. We do not include Maine, Missouri, or New Jersey because these states do not make email addresses of local election officials available. We do not include Maryland due to a clerical oversight.

We report other individual officials that were excluded from randomization, as well as reasons for these exclusions in Table A1. Local election officials were excluded from the study for concerns related to spillover, or multiple local election officials overseeing a single jurisdiction. All determinations were made prior to randomization. Figure A1 reports the Consort enrollment and randomization chart for this project.

Table A1: Local Election Officials excluded prior to randomization

Attrition by Study Exclusion Criteria

Exclusion Criteria Category	Exclusion Criteria Details	Number of deleted registrars or units of treatment (n)	Number of subjects remaining in cohort after exclusion (N)
Initial Count	Registrars from whom we collected public information		8104
Two levels of units per state	**County and municipality**		
	Delete registrars at county level - Wisconsin	(72)	8032
	Delete registrars at county level - Michigan	(83)	7949
	State and county		
	Delete registrars at state level - Delaware	(2)	7947
Missing emails	Delete registrars at county level with no email address - California, Idaho, Indiana, Maine, Missouri, Mississippi, New York, Pennsylvania	(652)	7295
	Delete registrars at municipality level with no email address - Connecticut, Michigan, New Hampshire, Rhode Island, Wisconsin	(183)	7112
Multiple registrars per unit of treatment	Randomly select one registrar per county and delete remaining duplicates:		
	Alabama	(3)	7109
	Arkansas	(19)	7090
	Connecticut	(79)	7011
	Louisiana	(15)	6996
	New Hampshire	(4)	6992
	Keep registrar with name and delete registrar with no name - Nevada	(2)	6990
	Keep registrar with job title "County Director" and delete registrar with job title "Deputy County" - Delaware	(6)	6984
	Keep registrar with job title "City Clerks" and delete registrars with job title "Town Clerks" - Michigan	(68)	6916
	For registrars with no job title, randomly select one and delete remaining duplicates - Michigan	(33)	6883
	Randomly select registrar based on ranking of job title (1- "city clerk", 2- "town clerk", 3- "village clerk"), delete remaining duplicates - Wisconsin	(230)	6653
Spillover - Registrars responsible of multiple units of treatment or registrars sharing email address	Randomly select one county, delete remaining counties for each registrar:		
	Georgia	(155)	6498
	Hawaii	(3)	6495
	Michigan	(31)	6464
	New York	(4)	6460
	South Dakota	(2)	6458
	West Virginia	(1)	6457
	Winsconsin	(7)	6450
Missing data	Unable to assign to treatment due to missing covariate data	(11)	6439
	Total	**(1665)**	**6,439**

Figure A1: CONSORT Document

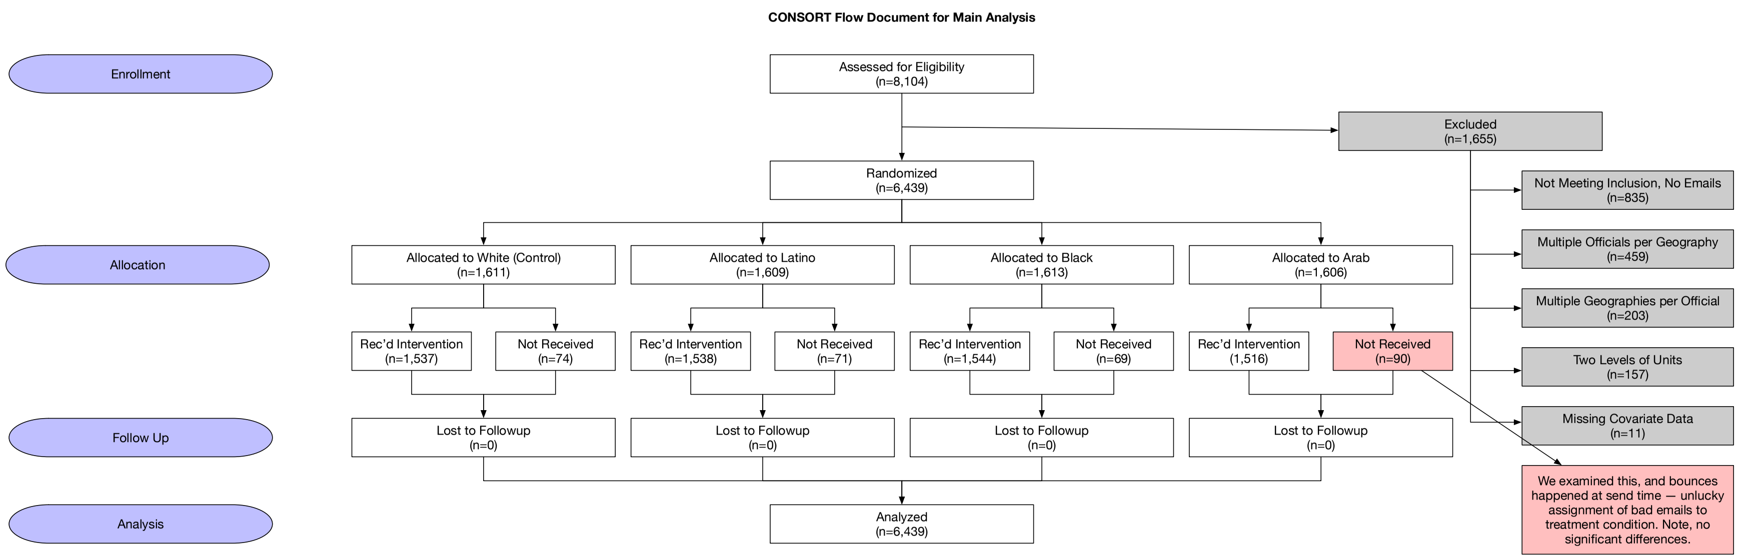

A2 Email Server Construction

At the design phase of the experiment, informed by the experience of White et al. (2015) we were concerned about the possibility that local elections officials might become aware of the conduct of our experiment.

A leading concern was that the domain name `ez-webmail` might structure election officials' responses. However, during the design phase of this experiment, we were surprised to make the observation that most client-side email services do not make the sender domain visible to the user. As we present in Figure A1, because we engineered our email server to match the `From:` name to the experimental stimulus, local election officials saw the sender name, not the email address in their Inbox. As a result, local election officials using most email programs would most likely not have seen the domain name of our sending server. We note, however that upon opening, all client-side programs make domain information visible to the election official (Figure A1).

Through the design of this experiment, our research into the front-end and back-end structure of how these emails are processed assuaged many of our concerns about the imperfect delivery of treatment. The following section describes this process.

A3 Email Back-End Considerations

Our leading concern was that these forms of contact would not reach election officials' inboxes. The primary cause of this failure is being captured by spam filters. To mitigate this concern, we expended significant IT effort to construct an email serving system that would be "well-respected" by client-side (i.e. election official side) email servers.

While the full technical specifications are beyond the scope of this article, we built the email sending server such that it was whitelisted for use on client-side email servers that included Gmail, Outlook, and Yahoo. We confirmed this, before sending, in two ways.

First, we use the *Return Path* sender score to evaluate that we had built sufficient sender score to have a high likelihood of reaching inboxes. Presently the leading indicator for delivering to client inboxes, the Return Path *senderscore* characterizes the reputation, and therefore probability of successful delivery of an email server. The identification of this product, as well as a considerable part of the sending architecture was influenced by interviews we conducted with leadership at a major direct-to-consumer (i.e. email) marketing firm.

Second, we tested that emails were actually arriving at inboxes. Specifically, we sent stimulus emails from our servers to a convenience sample of individuals associated with the research team, in an effort to cover a large part of the client-side landscape. We contacted colleagues, friends, and family using Microsoft Outlook at several different companies, and also contacted several people on each of Gmail and Yahoo email providers. These trials were instructive and serve as a cautionary tale for future researchers: in first rounds of pilot sending – trials where we had relatively low senderscore for our email server – we were not able to deliver any mail to any inbox.

Upon this realization, we took additional steps to improve the reputation of our email server. This involved server certificate signing, as well as ensuring that we had met spe-

(a)

(b)

(c)

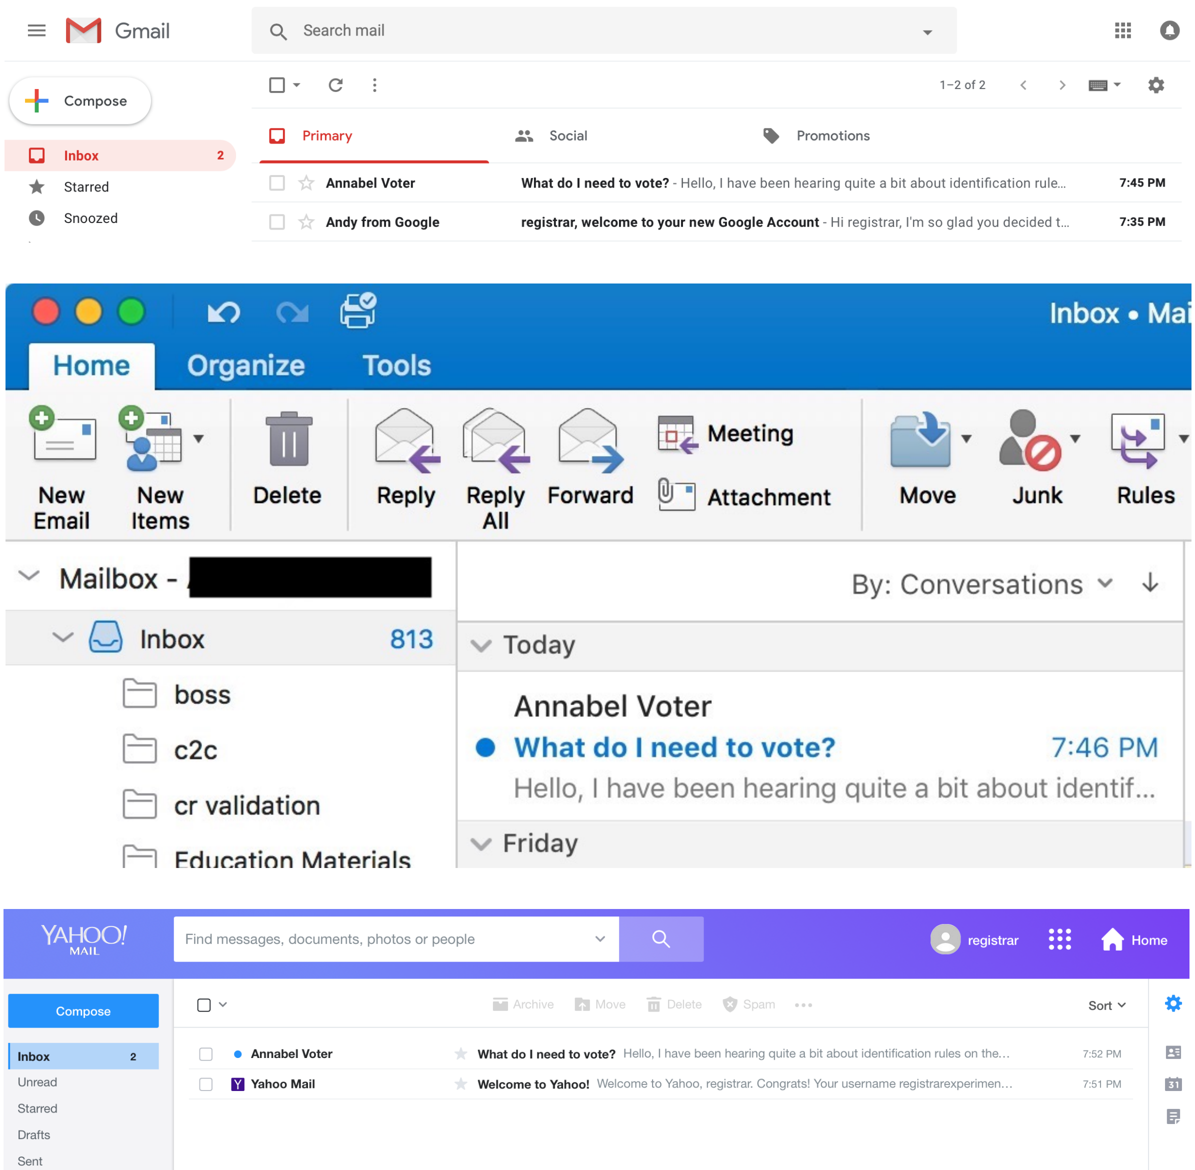

Figure A1: A likely view of our stimulus in local election officials' email inboxes. Subfigure (a) presents the view in gmail, (b) outlook, and (c) yahoo inboxes.

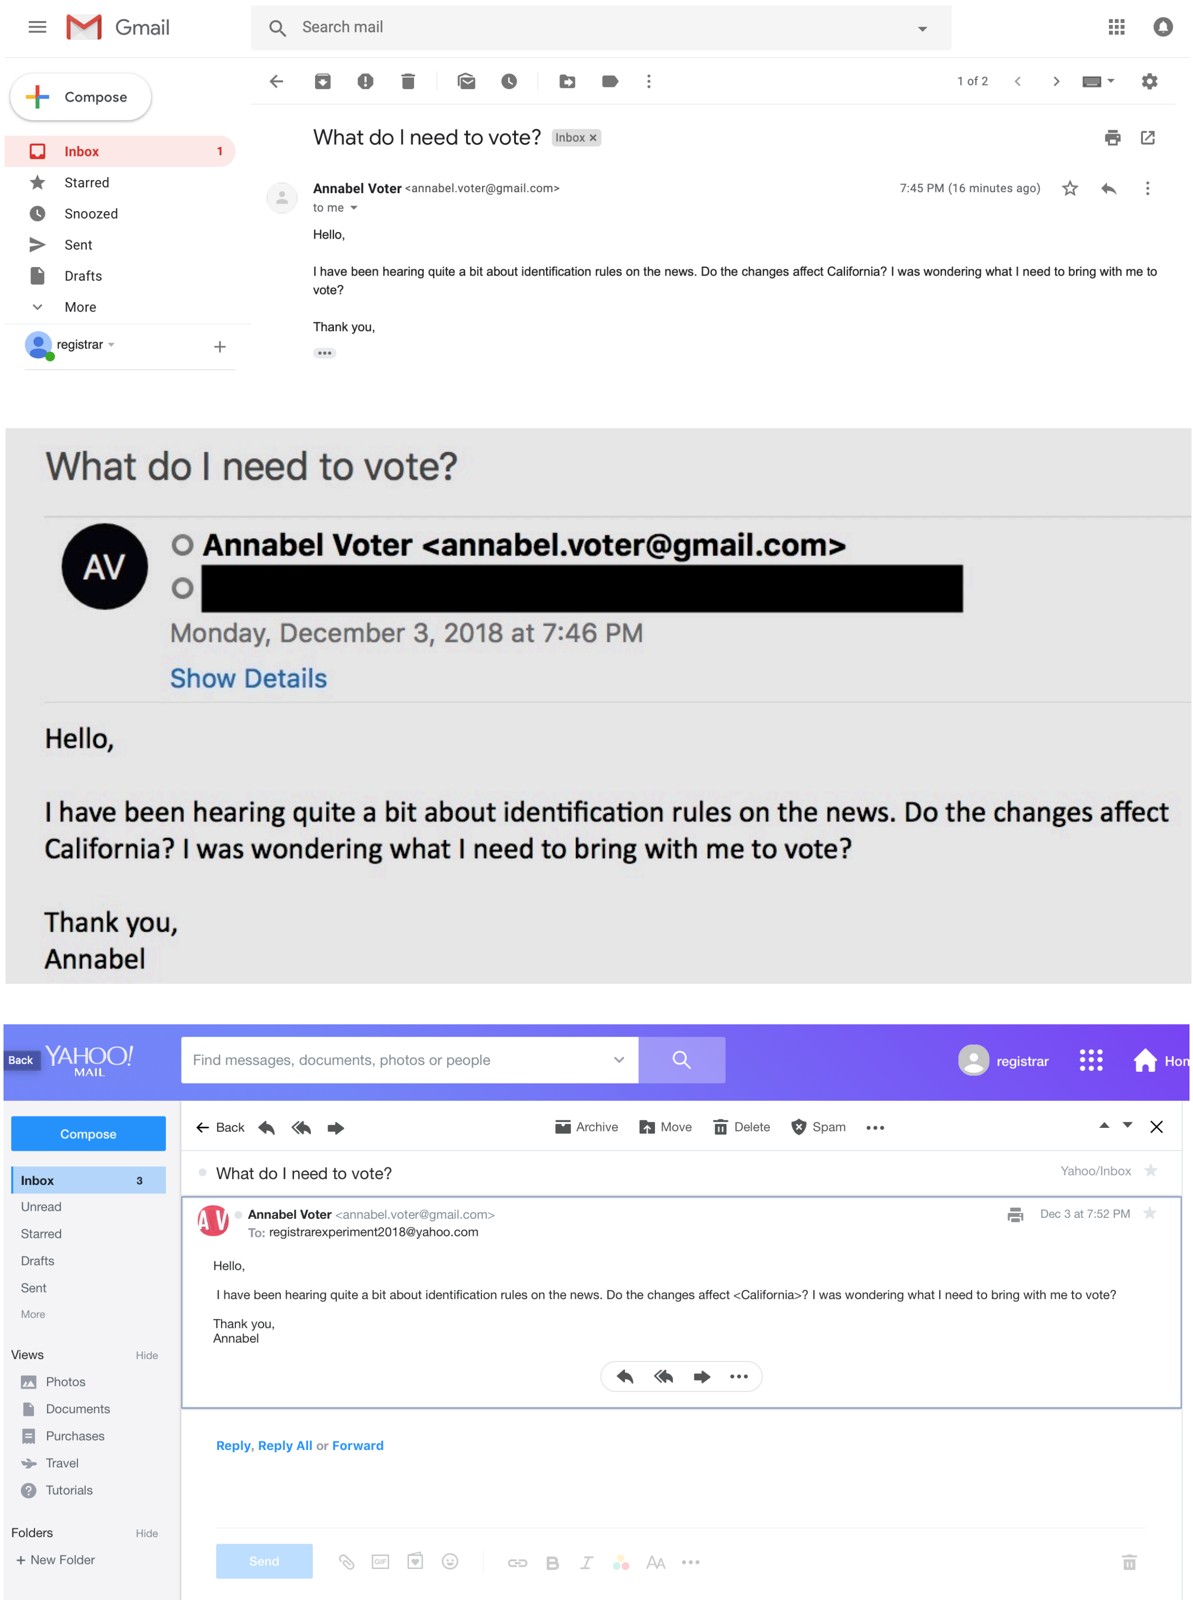

(a)

(b)

(c)

Figure A1: A likely view of our stimulus, once opened, in local election officials' email inboxes. Subfigure (a) presents the view in gmail, (b) outlook, and (c) yahoo inboxes.

cific (i.e. DKIM) authorization protocols. Although this is relatively routine for IT professionals, we would like to point out to experimentalists and those considering future audit studies that the process involved considerable work even for individuals with a background in this form of Information Technology.

Despite the cost and challenges of setting up a unique server, we would like to include this piece of dictum: It is our opinion that researchers who are going to engage in future audit studies should undertake the cost. On the one hand, the ability to flexibly define sender identity, include custom email headers and tracking infrastructure, and design custom data fields permits the estimation of theoretically interesting causal quantities (e.g. open rates). On the other hand, the increased cost of setting up this sending infrastructure serves to rebalance the costs bourn by experimenters and their audit/correspondence study subjects.

A4 Mailer Content

Unlike White et al. (2015), we did not vary whether the local election official receives a request directly related to voter identification. Because previous results establish that prejudicial behavior occurred almost exclusively in response to emails related to voter identification, we focus only on requests of that type.

To minimize the chance that local elections officials would become aware of the study, we took care to develop many versions of email language. In particular, all content that we mailed was a variant of a simple, three sentence paragraph that took the form: (1) Preamble; (2) Question One; (3) Question two.

By asking the same question in multiple ways, we achieve greater certainty that the resulting behavior is a response to the main causal variable of interest, the race of the putative voter, rather than any idiosyncratic feature of our request. Table A2 presents the different values for the preamble and the two questions. These elements were combined at random, to produce 27 variants of the message text delivered to local officials. These variants were scored by 171 humans for "*clarity*", "*warmth*" and "*appropriateness*". Data resulting from these evaluations suggest that the language variants would not be evaluated differently by readers.

As an example, one particular realization of our stimulus might draw the first cue each section, forming the email:

```
Dear <John Adams>,

I have been hearing quite a bit about identification rules on the
news. Do the changes affect <California>? I was wondering what I
need to bring with me to vote?

Thank you,

<Daniel Nash>
```

Cue Type	**Cue Text**
Preamble	I have been hearing quite a bit about identification rules on the news.
Preamble	I have heard a lot on the news about identification.
Preamble	The news has talked a lot about identification rules.
Question 1	Do the changes affect **state**?
Question 1	Are these changes happening in **state**?
Question 1	Do these affect **state**?
Question 2	I was wondering what I need to bring with me to vote?
Question 2	I was wondering if I need to bring anything specific with me to vote?
Question 2	Is there anything specific I need to bring to vote?

Table A2: Features manipulated for random assignment of messages to registrars of voters.

A5 Pilot

We conducted three pilots prior to deploying the experiment. The first pilot was conducted in Minnesota, chosen because it was the locale utilized as a pilot in previous studies White et al. (2015). Infrastructure problems meant that no emails were received by elections officials in the first pilot. We made changes, and conducted a second pilot in MN that successfully delivered emails. Finally, we conducted a third pilot in the western states of Washington, Oregon, California, and Nevada. These states were chosen due to their physical distance from other states, relatively small number of election officials, and peculiarities in election administration (e.g. Oregon does not conduct in-person elections).

A6 No Question Effects

In the following models, we report that the causal effects are invariant to including fixed effects for the specific questions asked.

Table A3

	Dependent variable:	
	GotResponse	
	(1)	(2)
Minority	-0.047^{***}	
	(0.014)	
Latino		-0.030^{*}
		(0.017)
Black		-0.0001
		(0.017)
Arab		-0.111^{***}
		(0.017)
Question Fixed Effect	Yes	Yes
Observations	6,439	6,439
R^2	0.006	0.013
Adjusted R^2	0.002	0.009
Residual Std. Error	0.493 (df = 6411)	0.492 (df = 6409)
Note:	$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$	

A7 Name Selection

In this appendix, we describe our approach to selecting the names of constituents. Our primary intent in choosing names from population lists was to eliminate the possibility of any name-based confounds to be responsible for differences in the behavior of local elections officials. Injecting variation in this facet of the treatment also lowers the likelihood that officials would become aware of the intervention by observing messages across offices sent from the same alias. By varying the names used to signal identity, we break from the general practice in political science, which has been to select a small number of names – frequently one or two for each racial/ethnic group (e.g., Butler and Broockman, 2011; White et al., 2015). Nevertheless our approach is in line with practices in the audit literature more broadly (see especially, Bertrand and Mullainathan, 2004)

In line with previous work on election official responsiveness, we exclusively use male names (White et al., 2015). Using names from a single gender reduces the variance in outcomes that is not associated with race or ethnicity signals, increasing the efficiency of the experimental design.

We draw white first names from the social security administration's records of births in Oregon in 1990. We utilize a list of distinctly African American names to produce our Black first names (Fryer and Levitt, 2004). Latino names are sourced from New York City baby names for children born between 2011 and 2014. Finally Arab/Muslim names were sourced from a list of common names (`http://www.behindthename.com/names/usage/arabic/`. Our intent in using this varied set of name sources was twofold. First and foremost, we wanted to generate plausible first names as an experimental stimulus. Second, we took care to ensure that the list of names we utilized was unlikely to match other name lists used in name-based audit studies.

We generate non-Hispanic White, Black, and Latino surnames from a US Census list of the 1000 most commonly occurring surnames (Word et al., 2008). This dataset provides information about the distribution of racial and ethnic groups by each surname. For example, among individuals with the most commonly occurring surname, Smith, the census data identifies that 73% identify as non-Hispanic White, 22% identify as Black, and 1.5% identify as Hispanic. To select names, we set minimum levels within each category. For a surname to be chosen as a white surname, more than 70% with that name needed claim a non-Hispanic White identity. For a surname to be chosen as a black surname, 30% or more of people with that surname needed claim a black identity; for Latino surnames we set this threshold at 60%. We note that this choice was made to produce what were, in our estimation, names that strongly signaled racial/ethnic group, without utilizing the *most* common surnames associated with these groups.

Arab/Muslim names, and indeed demographic and health statistics are difficult to identify (Al-Sayed et al., 2010). Consequently, we sourced surnames from `http://surnames.behindthename.com/names/usage/arabic`. This site does not provide frequency counts for names, so we assigned a uniform probability to each name being assigned.

With the set of first and last names created, we join the names together to produce a *given name* and *surname* pair that signals senders' racial/ethnic identities.

After constructing and curating a list of names to be sent as racial and ethnic primes, we recruited a set of workers through Amazon's *Mechanical Turk* (mTurk) worker platform. We paid mTurk workers a small amount to guess the probability that a particular name was of one or another ethnic group. Specifically, for each of 25 randomly selected names (from the set of $\approx$ 400) we asked workers to estimate their confidence (ranging from 0 percent to 100 percent) that an individual with a given name belonged to a particular racial or ethnic group.

As an example – the example we used in training workers for the mTurk task – we provided the name `Yao Ming`, a famous Chinese basketball player who played in the American NBA for 8 seasons. If a subject were certain that the name `Yao Ming` was a member of the Asian racial or ethnic group, the worker would place a certainty of 100 with this group. If the worker were mostly certain – for example 90 percent certain – that the name `Yao Ming` belonged to the Asian racial or ethnic group, she would place a 90 with that group and the remaining 10 percent certainty with other group(s) she thought the name may belong.

The results of this task are reported in section A17, Table A16.

A8 Blocking

We block on measures that are likely to predict whether a voting official will respond to (a) any form of contact and (b) forms of contact from minority voters. Specifically, we block on population density, proportion below 150 percent of the federal poverty line, proportion Black, proportion Latino, President Obama's margin of victory in the 2012 Presidential Election, and previous coverage by §5 of the VRA.

Our blocking data was most commonly measured at the county level – e.g. county electoral returns. However, the relevant electoral area addressed by a local election official may, or may not also be a county. In some states local election officials execute elections across multiple counties; in other states local elections officials represent a single county; while in still others officials might work at the municipal level. When our blocking features were more geographically broad than the area covered by a local election official, we apply the county level values to the municipal level. When our blocking features were more narrowly measured than the political geography covered by an official, we simply average the county-level measurements. Details of implementation can be found in the notebooks that accompany this work.

Blocking was implemented via the `blockTools` package written by Ryan Moore (Moore 2012.) Blocks of size four were created using an 'optimalGreedy' blocking algorithm. The algorithm begins by identifying the best pair of individual units to place in a single block, then identifies the best additional unit to include in that block, until the specified magnitude of the block is reached. It repeats the process until all units are blocked. We did not permit blocks from being formed between units in different states. In Table A4 we report the results of our blocking strategy. In brief, blocking and subsequent randomization succeeded.

Table A4

	ethnic_cue	Mean Density	Mean Income	Mean Black	Mean Latino	Mean Obama	Mean VRA
1	White	1.860	0.044	0.043	0.055	-0.063	0.120
		0.019	0.001	0.003	0.003	0.007	0.008
2	Latino	1.850	0.045	0.043	0.055	-0.061	0.117
		0.020	0.001	0.003	0.003	0.007	0.008
3	Black	1.850	0.045	0.044	0.056	-0.060	0.120
		0.020	0.001	0.003	0.003	0.007	0.008
4	Arab	1.840	0.045	0.043	0.054	-0.065	0.118
		0.020	0.001	0.003	0.003	0.007	0.008

Notes. Standard errors are reported beneath variable means

Table A5: Response Rates by Experimental Condition

Ethnic Cue	**White**	**Minority**	**Latino**	**Black**	**Arab**
Response Rate (%)	61.3	56.6	58.4	61.4	50.1
Standard Error	1.21	0.71	1.23	1.21	1.25
N	1,611	4,828	1,609	1,613	1,606

Notes: The *Minority* column includes all data from the *Latino*, *Black*, and *Arab* columns. Response rates and standard errors are reported in percentage terms.

A9 Nonparametric Results

The table reproduced in this section produces the non-parametric, difference in means between the white, minority, latino, black and Arab name-cues. As we report in Figure 1, minority, latino and Arab names receive responses at rates lower than white names. There is no detectable difference between the response rates of black and white names.

A10 Fixed Effects Models

Table A6 presents linear probability models estimating the same causality quantities reported in Figure 1 in the main body of the paper, though we provide more information in this Appendix. Models 1 and 2 estimate the causal effect of voter contact sent by non-white voters (model 1) and specific racial and ethnic classes of voters (model 2), but without including block-specific fixed effects. Models 3 and 4 estimate these same relationships, but include block fixed effects. Models 1 and 2 estimate robust (HC3) standard errors; models 3 and 4 estimate robust standard errors as constructed in the `lfe`, version *lfe_2.5-1998*.

We note that, while all models reported herein use *HC3* standard errors, we obtain substantively similar results when using Bell-McCaffery small-sample standard errors recommended by Lin and Green (2015).

In Model 1, we estimate that the local election officials respond to 61.3 percent of the emails they received from white voters. Emails received from racial and ethnic minority voters received a response at a rate 4.7 percent lower than this baseline: 56.6 percent of emails sent by minority names received a local election official response. Model 3 estimates the same relationship, but de-means the estimates within each block. The estimate of the causal relationship between sending an email as a minority voter rather than a white voter does not change substantively, although the blocking does improve the efficiency of the estimator.

In Models 2 and 4 we examine whether different racial and ethnic minority groups are treated differently by the local election officials. We find evidence to support this hypothesis. Models that do (Model 4) and do not (Model 2) include block fixed effects both find that emails from a Latino voter are 3.0 percent less likely to receive a response than emails sent from a white voter. In contrast, emails sent from Black voters are treated very similarly as emails sent from white voters. The estimate of the causal relationship is very nearly zero ($\beta = 0.1$ percent), and is roughly 1/30 the magnitude of the latino effect. In both Models 2 and 4 we estimate Arab/Muslim aliases receive a response from elections officials at a rate 11.3 percentage points lower than the baseline response rate.

Table A6: Causal Estimates

	GotResponse			
	(1)	(2)	(3)	(4)
Minority	-4.700^{***}		-4.710^{***}	
	(1.410)		(1.330)	
Latino		-2.970^{*}		-2.990^{*}
		(1.730)		(1.630)
Black		0.110		0.167
		(1.720)		(1.650)
Arab		-11.300^{***}		-11.300^{***}
		(1.740)		(1.630)
Constant	61.300^{***}	61.300^{***}		
	(1.210)	(1.210)		
Block FE	No	No	Yes	Yes
Observations	6,439	6,439	6,439	6,439
R^2	0.002	0.009	0.330	0.337
Note:		$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$		

A11 Robust to Link Function

While OLS estimators are unbiased estimates of the causal effect under this research design, we demonstrate that the choice of link function in a general linear model does not meaningfully alter estimates. In Table A7 and Table A8, we use a maximum likelihood approach to estimating these models, first with a gaussian link function, but also with logit and probit functions.

Table A7: Robust to Logit and Probit Specification

	Dependent variable:		
	GotResponse		
	normal	*logistic*	*probit*
	(1)	(2)	(3)
Minority	-0.047^{***}	-0.194^{***}	-0.121^{***}
	(0.014)	(0.059)	(0.037)
Intercept	0.613^{***}	0.461^{***}	0.288^{***}
	(0.012)	(0.051)	(0.032)
Observations	6,439	6,439	6,439
Log Likelihood	−4,589.000	−4,379.000	−4,379.000
Akaike Inf. Crit.	9,183.000	8,762.000	8,762.000
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table A8: Robust to Logit and Probit Specification

	Dependent variable:		
	GotResponse		
	normal	*logistic*	*probit*
	(1)	(2)	(3)
Latino	-0.030^{*}	-0.124^{*}	-0.077^{*}
	(0.017)	(0.072)	(0.045)
Black	0.001	0.005	0.003
	(0.017)	(0.072)	(0.045)
Arab	-0.113^{***}	-0.459^{***}	-0.286^{***}
	(0.017)	(0.072)	(0.045)
Intercept	0.613^{***}	0.461^{***}	0.288^{***}
	(0.012)	(0.051)	(0.032)
Observations	6,439	6,439	6,439
Log Likelihood	−4,567.000	−4,356.000	−4,356.000
Akaike Inf. Crit.	9,141.000	8,721.000	8,721.000
Note:			*p<0.1; **p<0.05; ***p<0.01

A12 Pilot Inclusion

We piloted our delivery and intake engineering in two separate pilots. The first, executed in Minnesota, was initially met with technical implementation issues – we received server information that no emails from our system were being delivered to local election official addresses. We addressed this issue, and, because our forensics determined that it would not be possible for officials to be aware of our first pilot, we re-ran this pilot and were successful on this follow-up attempt. To ensure that our engineering was not only a Minnesota-specific success, we ran a second pilot in the Western states of Washington, Oregon, California, and Nevada. We chose these states because of their relatively small local election official population (233 total local election officials), and their distance from locales with many local election officials.

As we report in Table A9 and Table A10, neither including nor excluding these pilot states from the analysis changes the substance or the interpretation of the core results. In addition, there is no evidence that the causal effect is different in pilot compared to non-pilot states.

Table A9: Robust to Pilot Exclusion

	Dependent variable:		
	GotResponse		
	(1)	(2)	(3)
Minority Cue	-0.047^{***}	-0.046^{***}	-0.046^{***}
	(0.014)	(0.014)	(0.014)
Pilot			0.120^{*}
			(0.065)
Minority Cue * Pilot			-0.034
			(0.076)
Constant	0.613^{***}	0.609^{***}	0.609^{***}
	(0.012)	(0.013)	(0.013)
Include Pilot	Yes	No	Yes
Observations	6,439	6,206	6,439
R^2	0.002	0.002	0.003
Adjusted R^2	0.002	0.001	0.003
Note:	*p<0.1; **p<0.05; ***p<0.01		

Table A10: Robust to Pilot Exclusion

	Dependent variable:		
	GotResponse		
	(1)	(2)	(3)
Latino Cue	-0.030^{*}	-0.030^{*}	-0.030^{*}
	(0.017)	(0.018)	(0.018)
Black Cue	0.001	0.005	0.005
	(0.017)	(0.018)	(0.018)
Arab Cue	-0.113^{***}	-0.112^{***}	-0.112^{***}
	(0.017)	(0.018)	(0.018)
Pilot			0.120^{*}
			(0.065)
Latino Cue * Pilot			0.021
			(0.093)
Black Cue * Pilot			-0.107
			(0.092)
Arab Cue * Pilot			-0.013
			(0.093)
Constant	0.613^{***}	0.609^{***}	0.609^{***}
	(0.012)	(0.012)	(0.012)
Include Pilot	Yes	No	Yes
Observations	6,439	6,206	6,439
R^2	0.009	0.009	0.010
Adjusted R^2	0.008	0.009	0.009
Note:			$^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

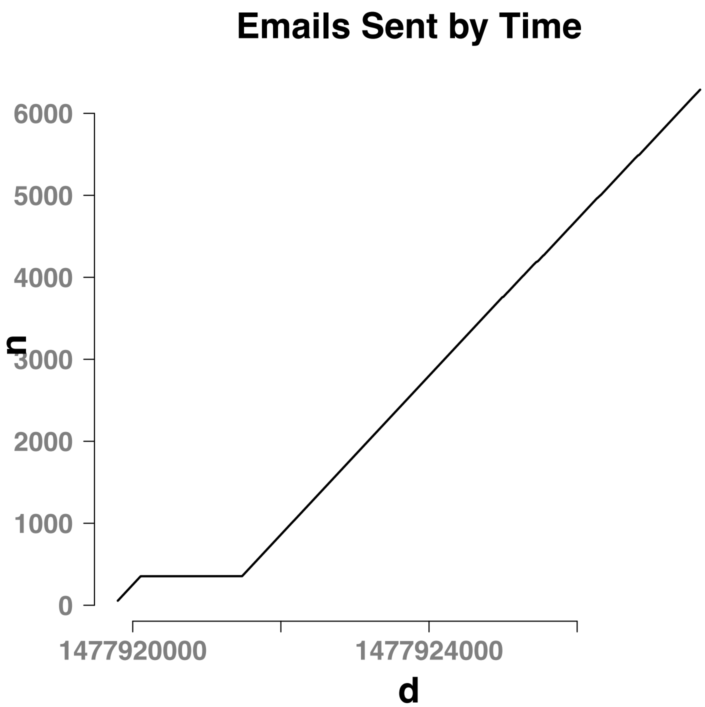

h

Figure A10: The number of emails sent is marked on the y-axis, and the time (in UNIX seconds, in the UNIX epoch) are plotted on the x-axis. Note the 30 minute gap in sending. Here, we waited to ensure that emails were making it to officials' inboxes, before green-lighting the remainder of the production email run.

A13 Email Send Timing

In this appendix, we describe the timing of sending our emails. Emails were delivered in waves over a few hours to officials in the sample. We decided against emailing all local election officials at the same time to reduce the chance of unexpected results due to technical errors and to reduce possible spillover effects. We also considered emailing local election officials over a period of multiple days. Ultimately, we were concerned that the likelihood of differential response rates on different days outweighed the benefits to spreading email messages across several days. Note the 30 minute gap in sending. Here, we waited to ensure that emails were making it to officials' inboxes, before green-lighting the remainder of the production email run. We determined that our stimulus was making it to election officials inboxes when we received replies from officials in several states.

A14 Time to Response

In this appendix, we consider how much time was required for local election officials to respond to our email. To do so, we merge tracker hits from our server with the time that we received an email reply. The tracker hit records when a registar opened the email, and the response effectively records when the task is complete.

We take some care in computing this, because election-official-side email clients handle our tracker hits differently. In particular, some email clients "cache" a version of our image on their own servers to speed up the loading of images in emails. When this occurs, we do not receive reliable information about when an email was opened.

We work around this problem by including only the *first* load that occurs on our sever. Not only does this preclude problems with individuals' email clients, but at the same time we believe it also represents a conservative (long) estimate of the time to complete the task.

As we plot in Figure A10, the task that we set before election officials did not require a substantial amount of time. Of those responses that we received, and have valid data for, the median time to respond was fewer than three minutes. It is, however, important to note that we neither have information about the time to respond for officials who do not respond to our stimulus, nor for officials whose email clients prohibit us from gathering reliable data.

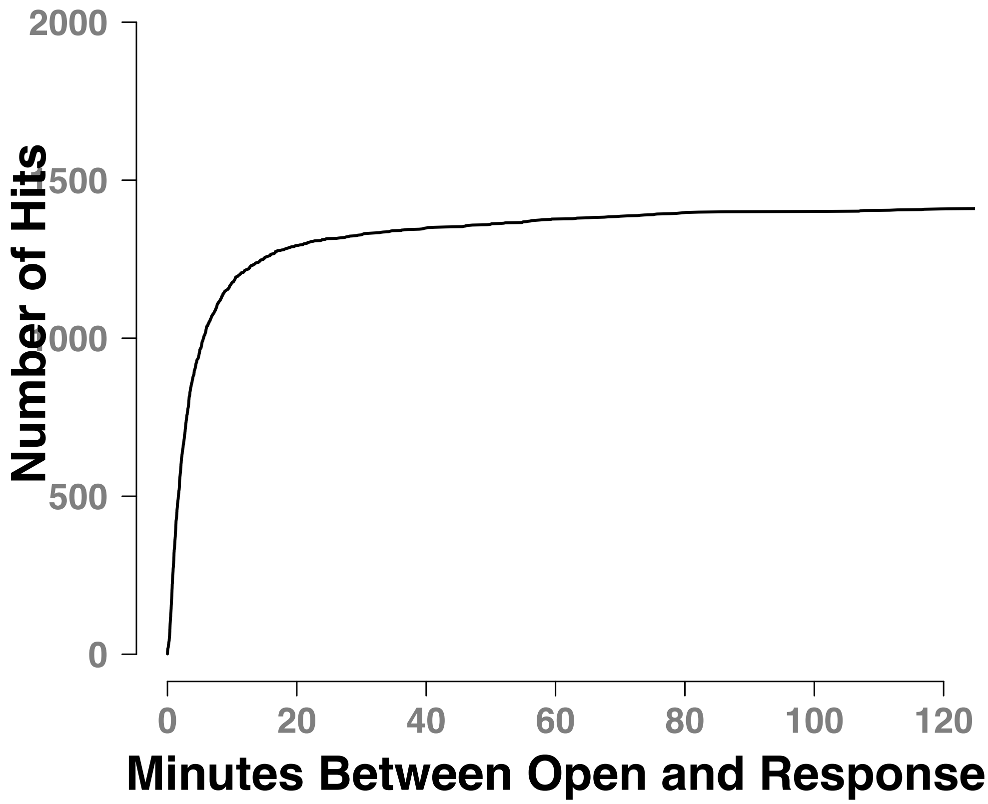

Figure A10: On the x-axis are the minutes elapsed since the first time the local election officials opened our stimulus, until the time that we received a response from that election official. On the y-axis are the cumulative number of responses that have been received in that duration of time.

A15 No Damage from Spillover

After we collected outcome data, we learned that election officials in some states were suspicious about the emails, and contacted their state organization who, in turn, contacted the national organization. As well, we came to learn that at least one other research team was pursuing a substantively similar project, using the domain registered by White et al. (2015).

To examine whether this notification seems to have affected the willingness of elections officials to respond we estimate two distinct robustness checks. First we estimate a number of Cox proportional hazard (duration) models. We choose this model class because they are unbiased in the presence of censored data. In particular, this model type permits us to estimate models that use the pre-registered end date of observation, as well as the timing of the NASS clerk email as the end date of observation. As we report in Table A11, the coefficients estimated in all models are highly stable.

As a second robustness check, we estimate our core, pre-registered models again, but excluding states where the news reported early awareness: Michigan, New Hampshire, and Colorado. The results we report in Table A12 retain their statistical significance and substantive interpretation. Although these are not dispositive tests, this set of results do not surface any evidence to suggest that the differences in response rates we observe are being caused by awareness.

Table A11: Cox Proportional Hazards Models

	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Minority Cue	-0.13^{***}	-0.14^{***}	-0.13^{***}	-0.13^{***}				
	(0.04)	(0.04)	(0.04)	(0.04)				
Latino Cue					-0.10^{*}	-0.10^{*}	-0.08^{*}	-0.07
					(0.05)	(0.05)	(0.05)	(0.05)
Black Cue					-0.02	-0.03	-0.01	-0.02
					(0.05)	(0.05)	(0.04)	(0.04)
Arab Cue					-0.29^{***}	-0.29^{***}	-0.31^{***}	-0.30^{***}
					(0.05)	(0.05)	(0.05)	(0.05)
Data Subset	Clean	Clean	All	All	Clean	Clean	All	All
Censoring Date	Election	Clerk	Election	Clerk	Election	Clerk	Election	Clerk
Observations	4,548	4,548	6,435	6,435	4,548	4,548	6,435	6,435
R^2	0.002	0.002	0.002	0.002	0.01	0.01	0.01	0.01

Notes. Cox proportional hazards models. Outcome is converting from no response to response. *Clean* data subset are states without known spillover, and exclude pilot data. *All* data subset includes all states' data. Two censoring points are estimted. *Election* is the pre-registered censoring date at election day; *Clerk* places the censoring date at the time of the NASS email notification. $^{*}p<0.1$; $^{**}p<0.05$; $^{***}p<0.01$

Table A12: No Difference in Estimates in Interference States

	Dependent variable:							
	GotResponse				HitTracker			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Minority	-0.045^{***}	-0.048^{***}			-0.033^{**}	-0.036^{**}		
	(0.013)	(0.016)			(0.014)	(0.016)		
Latino			-0.028^{*}	-0.036^{*}			-0.005	-0.001
			(0.016)	(0.019)			(0.017)	(0.020)
Black			0.004	-0.002			-0.006	-0.016
			(0.016)	(0.019)			(0.017)	(0.020)
Arab			-0.110^{***}	-0.107^{***}			-0.088^{***}	-0.092^{***}
			(0.016)	(0.019)			(0.017)	(0.020)
ln(pop dens)	0.152^{***}	0.117^{***}	0.149^{***}	0.112^{***}	0.063^{**}	0.059^{*}	0.060^{**}	0.054
	(0.029)	(0.035)	(0.029)	(0.034)	(0.029)	(0.035)	(0.029)	(0.035)
Pct < 150	(0.693)	(0.766)	(0.689)	(0.763)	(0.700)	(0.785)	(0.698)	(0.783)
Pct Black	-0.498^{**}	-0.227	-0.483^{**}	-0.213	-0.033	0.508^{*}	-0.014	0.532^{*}
	(0.236)	(0.301)	(0.235)	(0.300)	(0.239)	(0.309)	(0.238)	(0.308)
Pct Latino	-0.355^{*}	-0.161	-0.367^{*}	-0.179	-0.333	-0.177	-0.341^{*}	-0.192
	(0.202)	(0.241)	(0.201)	(0.240)	(0.204)	(0.247)	(0.203)	(0.246)
Obama Margin	-0.006	0.058	-0.015	0.053	-0.031	0.006	-0.039	0.001
	(0.084)	(0.095)	(0.083)	(0.095)	(0.085)	(0.098)	(0.084)	(0.097)
Observations	6,439	4,552	6,439	4,552	6,439	4,552	6,439	4,552
R^2	0.334	0.327	0.341	0.332	0.282	0.284	0.287	0.289
Adjusted R^2	0.109	0.097	0.118	0.104	0.039	0.040	0.046	0.046

Note: *p<0.1; **p<0.05; ***p<0.01

A16 Limited District Characteristic Heterogeneity

In the following models, reported in Table A13, Table A14, Table A15, we examine whether officials' response to treatment is different conditional on characteristics of their district. In particular, one hypothesis is that officials who preside over jurisdictions that hold a relatively large share of minority voters may be more likely to respond to a question about voting from minority voters. Indeed, as we show in Table A13 and Table A14, while there is little change in the responsiveness of election officials as the proportion of voters in that jurisdiction becomes increasingly black (shown in *Model (2)* and *Model (3)* in both Table A13 and Table A14), as we report in *Model (1)* in Table A13 and Table A14, there is some evidence that officials' responsiveness changes as the proportion of Latinos in a jurisdiction increases.

Of particular interest is the possibility that the large treatment effects for the Arab/-Muslim cue are driven by the implausibility of the treatment, due to the very small proportion of Arab Americans living in many jurisdictions. The results below are motivated by the following logic: if treatment effects for a given identity are driven by implausibility then they should be smaller in places where individuals who have been ascribed that identity are more numerous.

The distribution of Arab Americans is somewhat distinct from the distribution of blacks and Latinos. Indeed, data from the current CPS suggests that just 8 percent of U.S. counties have no Latino population, and 25 percent have no black population. In contrast, fully half of the counties in the U.S. have no residents who identify with an Arab heritage. Thus, it is possible that the lack of variation in the `pct_arab` population variable has made it mechanically impossible for a regression to detect a heterogeneous treatment effect.

To examine whether this is possible, we rescale the percent of Arab population into a three-level factor variable in the following way:

- For geographies that have zero Arab population, we code the rescaled variable as 0. This represents the 0-50th percentile distribution of communities arranged by Arab-American population;
- Among geographies that have at least one person who identified an Arab heritage, we make a further split at the median.
  - The lower of the two groups, the set of communities that represent the 50-75th percentile distribution of communities; and,
  - The higher of the two groups, the set of communities that represent the 75-100th percentile distribution of communities.

As noted, this indicator splits the Arab population into three categories. The first category covers the 50 percent of U.S. counties with no Arab population. The second covers the 25 percent of U.S. counties whose Arab-American population is below the median value for those counties in which any Arabs live. In these counties, Arab Americans still represent a small part of the population: 0.12%. The third category covers the remaining 25 percent of counties whose Arab population is above this median. In these counties

Table A13

	Dependent variable:		
	GotResponse		
	(1)	(2)	(3)
Minority	-0.052^{***}	-0.048^{***}	-0.044^{***}
	(0.015)	(0.015)	(0.015)
Percent Latino	-0.241		
	(0.236)		
Percent Latino × Minority	0.093		
	(0.143)		
Percent Black		-0.163	
		(0.230)	
Percent Black × Minority		0.013	
		(0.133)	
Percent Arab			1.580
			(2.440)
Percent Arab × Minority			-1.270
			(2.530)
Observations	6,439	6,439	6,406
R^2	0.330	0.330	0.329
Adjusted R^2	0.104	0.103	0.101
Note:			*p<0.1; **p<0.05; ***p<0.01

Table A14

	Dependent variable:		
	GotResponse		
	(1)	(2)	(3)
Latino	-0.049^{***}	-0.026	-0.028
	(0.019)	(0.018)	(0.018)
Black	0.013	-0.003	0.003
	(0.019)	(0.018)	(0.018)
Arab	-0.121^{***}	-0.113^{***}	-0.109^{***}
	(0.019)	(0.018)	(0.018)
Percent Latino	-0.227		
	(0.233)		
Percent Latino × Latino	0.345^{**}		
	(0.167)		
Percent Latino × Black	-0.199		
	(0.174)		
Percent Latino × Arab	0.138		
	(0.168)		
Percent Black		-0.173	
		(0.234)	
Percent Black × Latino		-0.098	
		(0.162)	
Percent Black × Black		0.119	
		(0.166)	
Percent Black × Arab		0.008	
		(0.156)	
Percent Arab			1.680
			(2.460)
Percent Arab × Latino			-0.850
			(2.780)
Percent Arab × Black			-0.657
			(2.770)
Percent Arab × Arab			-1.740
			(2.670)
Block FE	*Yes*	*Yes*	*Yes*
Observations	6,439	6,439	6,406
R^2	0.339	0.337	0.337
Note:			*p<0.1; **p<0.05; ***p<0.01

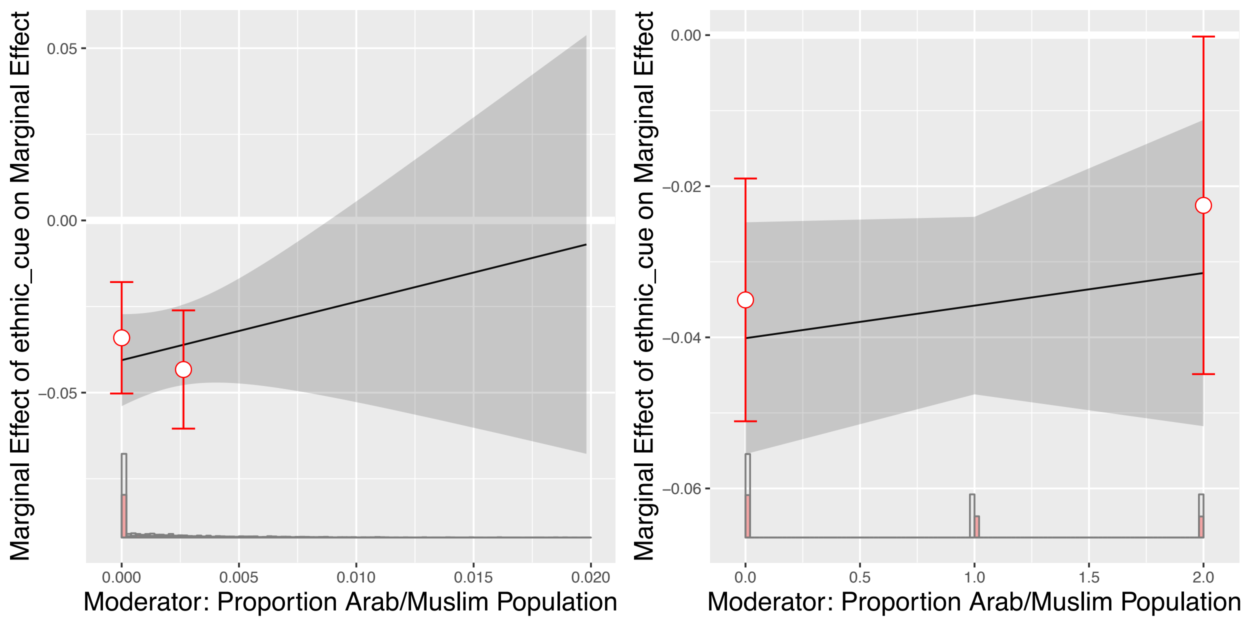

Figure A15: Flexible estimates of HTE across Arab/Muslim population (left plot) and the three-level factor that indicates (0.0) zero arab/mulsim population; (1.0) Less than 1% Arab/Muslim population; and (2.0) 1% or more Arab/Muslim population.

with the greatest presence of Arab-Americans, this group represents, on average, 1% of the county population. In the geographies corresponding to this quartile of the distribution, it would not be uncommon for a local election official to be in the presence of an Arab-American person or family at a community gathering of several hundred people – such as a parade or high school graduation.

As we report in Table A15 after rescaling the data in this way, there is little evidence that our treatment effects were moderated in geographies in which census data records a greater number of Arab Americans. Neither of the terms interacting the treatment with the recoded covariate described above yield point estimates with p-values approaching standard thresholds of statistical significance. (We recognize that this failure to reject could be driven by insufficient statistical power.)

To provide further evidence, Figure A15 reports estimates of the treatment effect of receiving an email from an Arab/Muslim sender rather than a white sender, using the `iterflex` method that flexibly estimates and projects treatment effects across moderating variables (Hainmueller et al., 2018). We note in the left plot that the uncertainty estimates rapidly expand among counties with larger arab populations due to the sparse nature of the data: there are only 5 counties with an Arab/Muslim population larger than 10%. On balance, the evidence presented here conforms to the logic presented above in support of our argument: treatments were not more influential in those places where Arab Americans are less numerous.

Table A15

	Dependent variable:	
	GotResponse	
	(1)	(2)
Minority Cue	-0.040^{**}	
	(0.021)	
Latino Cue		-0.020
		(0.025)
Black Cue		0.004
		(0.026)
Arab Cue		-0.106^{***}
		(0.025)
1-50pct Arab	0.073^{**}	0.073^{**}
	(0.032)	(0.032)
51-100pct Arab	0.082^{**}	0.081^{**}
	(0.034)	(0.034)
Minority Cue * 1-50pct Arab	-0.026	
	(0.035)	
Minority Cue * 51-100pct Arab	0.005	
	(0.036)	
Latino Cue * 1-50pct Arab		-0.024
		(0.043)
Black Cue * 1-50pct Arab		-0.004
		(0.043)
Arab Cue * 1-50pct Arab		-0.050
		(0.043)
Latino Cue * 51-100pct Arab		-0.008
		(0.044)
Black Cue * 51-100pct Arab		-0.0004
		(0.044)
Arab Cue * 51-100pct Arab		0.026
		(0.044)
Block FE	*Yes*	*Yes*
Observations	6,439	6,439
R^2	0.332	0.340
Adjusted R^2	0.107	0.115
Note:	*p<0.1; **p<0.05; ***p<0.01	

A17 Names and Assessment of Racial and Ethnic Group

Table A16: Name Score Table

Name	Ethnic Cue	Mean White	Mean Latino	Mean Black	Mean Arab
Daniel Nash	White	97.6	0.9	1	0
Mathew Roberts	White	95	0	3.7	0
Alex Steele	White	94.6	0.4	5	0
Nicholas Austin	White	94.6	0.4	4.6	0
Zachary Fitzpatrick	White	94.3	0.7	4.1	0
Christopher Schmidt	White	93.7	0.1	3.4	0.1
Ryan Thompson	White	93.1	0	6.2	0
Timothy Bartlett	White	93	0	6	0
Corey Kennedy	White	93	0	7	0
Garrett Riddle	White	92.9	0.4	6.6	0
Austin Walsh	White	92.4	0.3	5.8	0
Christopher Rogers	White	92.1	0	7.9	0
Jacob Gates	White	92	0	6.7	0
Kyle Caldwell	White	92	0	6	0
Matthew Pratt	White	91.4	0	8.6	0
Joseph Mayer	White	91.3	0	8.7	0
Ian Thornton	White	90.5	0	9.5	0
Scott Sherman	White	89.5	0.2	8.8	0
Daniel Horn	White	89.3	0	2.5	0
Zachary Proctor	White	89	0	7.5	0
Brandon Hart	White	88.8	0	11.2	0
Nathan Brewer	White	88.3	0	2.8	0
Garrett Allen	White	87.5	0.6	11.9	0
John Miller	White	87.3	0	10.9	0
Robert Peterson	White	87.2	0	11.7	0
Dylan Garrett	White	86.9	0	7.5	0
Michael Quinn	White	86.7	0	13.3	0
Justin Kramer	White	86.4	0	8.2	0
Robert Todd	White	86.1	0.4	12.1	0
Travis Roberts	White	85.7	0.7	10.7	0
Richard Bowers	White	85.7	1.3	6.7	0
Jason Gillespie	White	85.4	0.4	7.1	0
Garrett Miller	White	85.3	0	14.7	0
Kyle Thompson	White	84.4	0	15	0
Dustin Lawson	White	84.2	0	15.3	0
Sean Cooper	White	84.1	0	15.3	0
James McPherson	White	83.2	0	14.6	0
Brandon Pierce	White	83.2	0.5	14.7	0
John Gregory	White	83	2.9	10.2	0
David Cochran	White	82.9	0	17.1	0
Seth Rodgers	White	82.9	0.7	6.4	1.4
Christopher Anderson	White	82.9	0.2	16.8	0
Tyler Reeves	White	82.5	0.4	12.9	0
Justin McIntyre	White	82.5	5.6	6.4	0
Matthew Moore	White	82.4	0.7	16.6	0.1
Stephen Peterson	White	81.9	0	16.2	0
Kyle French	White	81.8	0.9	13.6	0
Timothy Middleton	White	81.4	0	17.7	0

Ian Smith	White	81.3	0	18.7	0
Tyler Larson	White	81.1	0	18.9	0
Gregory Leblanc	White	80.8	0.4	11.5	1.5
Ryan Chapman	White	80.7	0.2	16.8	0
William Humphrey	White	80.6	0	19.4	0
Justin Mullins	White	80.5	0	11.4	0
Joshua Burke	White	80.4	0	14.2	0
Jacob Haas	White	80	0	2.2	0
Levi Wolfe	White	80	0	0	0
Kevin Patterson	White	80	0	19.1	0
Jeremy Short	White	79.6	0	18.7	0
Cody Lang	White	79.4	0	3.1	0
Taylor Long	White	79	0	17.7	0
Zachary Bailey	White	78.8	0	12	0
Michael White	White	77.8	0	16.7	0
Jeffrey Phillips	White	77.1	0.4	21.7	0
Travis Miller	White	77.0	0	23.0	0
Brian Bennett	White	76.9	0	19.4	1.2
Robert Cochran	White	76.4	2.3	12.7	4.5
Michael Hendrix	White	76.2	0	17.9	0
Travis Osborn	White	75.4	0.8	7.1	0
Michael Boyer	White	75.3	0	15.3	1.3
Travis Collins	White	75	0	24.3	0
Christopher Hebert	White	74.7	0.7	22.7	0
Samuel Peters	White	74.5	0	18.2	0
Shane Page	White	74.4	1.2	24.4	0
Jeffrey Fox	White	74.4	0.8	8.1	0
Anthony Underwood	White	73.8	0	23.8	0
Justin Lyons	White	73.5	6.7	18.0	0
Michael Rose	White	71.9	3.8	23.1	0
Devin Foster	White	71	0	27	0
Joshua Clark	White	70	0	5	0
Jordan Rogers	White	69.7	0	21.6	0
Joseph Graves	White	68.8	0	17.8	6.2
Robert Reed	White	68.2	1.7	10.2	16.7
Tyler Murray	White	67.3	2	24	1.3
James Marsh	White	66.9	1.2	13.8	0
Travis Frye	White	66.8	0	24.1	0
Cameron Young	White	65.6	0	23.7	0
Stephen Sherman	White	64.6	0	26.9	0
Benjamin Wood	White	64	0	14.5	0
Eric Murray	White	61	0	29	0
Andrew Allen	White	60.9	0	28.4	0
Austin Hall	White	59.5	0	24.1	1.8
Samuel Wood	White	55.8	0	44.2	0
Marcus McFarland	White	55.5	0	44.5	0
Michael Lang	White	55.5	2.7	12.3	0
Samuel Hopkins	White	51.2	0	34.6	1.7
Brandon Estes	White	50.8	36.6	11.6	0
Sean Watts	White	40.4	1.8	50.7	1.4
Jordan Smith	White	39.6	0	50.4	0
Jose Hanson	White	9.5	77.5	12.5	0
Jose Cruz	Latino	0	100	0	0
Jorge Castro	Latino	0	100	0	0

Cesar Marquez	Latino	0	100	0	0
Jose Gutierrez	Latino	0	100	0	0
Juan Campos	Latino	0	100	0	0
Saul Gonzalez	Latino	0	100	0	0
Miguel Salazar	Latino	0	100	0	0
Jesus Perez	Latino	0	100	0	0
Diego Velazquez	Latino	0	100	0	0
Fernando Hernandez	Latino	0	100	0	0
Juan Ramos	Latino	0	99.6	0	0
Jose Valdez	Latino	0	99.6	0.4	0
Edwin Vasquez	Latino	0.6	99.4	0	0
Gerardo Escobar	Latino	0.8	99.2	0	0
Esteban Herrera	Latino	0	99.2	0	0
Jose Mendez	Latino	0	98.2	0.7	0
Luis Gomez	Latino	1.1	97.9	0.5	0
Fernando Acosta	Latino	1.1	97.8	0	0
Adriel Hernandez	Latino	0.8	97.3	1.2	0
Aldo Garcia	Latino	0	97.3	0	0
Jaime Gonzalez	Latino	1.4	97.1	1.4	0
Alejandro Rodriguez	Latino	0	96.9	3.1	0
Emilio Gonzalez	Latino	0.4	96.8	2.1	0
Esteban Contreras	Latino	2.3	96.6	0	0
Dariel Valdez	Latino	0	96.2	1.2	0
Enrique Lopez	Latino	3.8	96.2	0	0
Camilo Lopez	Latino	1.1	96.1	0	0
Miguel Barrera	Latino	0.7	95.7	1.8	0
Angel Ruiz	Latino	2	95.5	0.5	0
Roberto Reyes	Latino	0	95	5	0
Edwin Santiago	Latino	5.4	94.6	0	0
Angel Navarro	Latino	0	94.4	5.6	0
Ricardo Gomez	Latino	0.7	94.3	0.3	0
Marvin Lopez	Latino	3.6	92.7	2.7	0
Alejandro Ibarra	Latino	0.4	92.7	2.7	0
Jesus Hernandez	Latino	1.3	92.3	1.7	1.3
Emilio Cabrera	Latino	7.7	92.3	0	0
Cristian Ramirez	Latino	1.2	92.2	0	0
Jesus Martinez	Latino	2.1	92.1	1.4	1.4
Julio Morales	Latino	0.4	92.1	0	7.1
Adan Perez	Latino	2.5	91.5	0	0
Angel Maldonado	Latino	3.8	91.2	0	0
Darwin Gonzales	Latino	4.2	90.8	4.6	0
Dariel Garcia	Latino	2.1	90.7	6.4	0
Esteban Jimenez	Latino	0	90.4	1.9	0
Alberto Mendoza	Latino	0.7	90	1.4	0
Edgar Garcia	Latino	9	90	1	0
Miguel Rubio	Latino	0	89.1	9.1	0
Pablo Escobar	Latino	5.6	88.9	0	5.6
Luis Martinez	Latino	0	88.9	11.1	0
Carlos Villarreal	Latino	1.9	88.8	0.8	0
Luis Gonzalez	Latino	3.3	88.3	0	0
Jean Lopez	Latino	7.9	88.2	2.6	0
Carlos Ramos	Latino	1.4	88.2	0	0
Juan Perez	Latino	2.5	86.7	10.8	0
Ricardo Garza	Latino	5.8	86.7	1.7	1.7

Manuel Padilla	Latino	0	86.4	0	4.3
Miguel Rodriguez	Latino	1.8	86.4	0.9	0
Angel Pineda	Latino	5	85	1.2	1.2
Luis Moreno	Latino	2.5	84.6	0	0
Iker Martinez	Latino	3.2	83.9	1.1	0.7
Edgar Cardenas	Latino	8.7	83.7	1.7	0
Edwin Hernandez	Latino	11.1	83.5	3	0.5
Mario Chavez	Latino	3.6	82.1	1.4	1.4
Johan Estrada	Latino	8.3	80.7	0.9	0.7
Jefferson Sanchez	Latino	9.3	80.7	9.3	0
Johan Garcia	Latino	11.7	80.6	3.9	0
Emiliano Lopez	Latino	1.7	80	1.7	1.7
Erick Hernandez	Latino	13.8	79.4	5.3	0
Giovani Herrera	Latino	14.2	79.2	0	1.7
Luis Padilla	Latino	3.5	78.8	1.9	0
Randy Munoz	Latino	14.5	78.8	0	0
Jadiel Rodriguez	Latino	1.7	78.8	15.8	0.4
Brayan Estrada	Latino	2.8	78.2	9.5	1
Erik Rodriguez	Latino	7.7	78.2	0.5	0
Erick Suarez	Latino	13.5	76.9	2.7	1.5
Maximo Flores	Latino	9.7	76.1	3.2	0
Yaniel Campos	Latino	1.2	74.4	5.9	1.2
Miguel Trevino	Latino	0.9	72.6	5	0
Yair Fuentes	Latino	0	69.5	4.1	18.2
Matias Murillo	Latino	4.8	69	1	6
Anderson Guerrero	Latino	18.8	68.8	2.5	1.2
Edwin Castaneda	Latino	21.1	68.2	0	0
Kenny Rodriguez	Latino	27.1	67.4	0.9	1.2
Damian Martinez	Latino	13.7	66.8	18.2	0
Januel Aguilar	Latino	7.2	66.1	8.3	1.7
Noel Torres	Latino	22.3	65.9	11.8	0
Ismael Romero	Latino	5.8	60.4	4.2	24.6
Derick Torres	Latino	21.8	59.5	13.2	1.8
Julius Salazar	Latino	8.8	58.4	2.2	8.8
Angel Ponce	Latino	14.2	52.8	19.2	1.1
Thiago Zamora	Latino	2	52.5	6.5	6
Junior Delgado	Latino	15	50.4	30	0
Kenny Lozano	Latino	35.4	45.7	8.9	0
Jael Calderon	Latino	13.3	44	29.3	0
Darwin Guzman	Latino	26.0	42.4	17.4	0.7
Edwin Zuniga	Latino	12.7	38.7	22.7	3.3
Byron Salazar	Latino	34.2	31.5	24.6	6.9
Jean Barrera	Latino	45	23	5	2
Jefferson Ponce	Latino	55.9	0.5	28.2	0
DeShawn Jackson	Black	2.4	0	97.6	0
Tyrone Brown	Black	1.2	1.7	96.7	0
DeShawn Harris	Black	2.9	0.3	96.7	0
DeShawn Brown	Black	2.1	0	96.7	0
Darius Thomas	Black	2.5	0	96.2	1.2
DeAndre Jackson	Black	1.4	0.8	96.1	0
Jamal Jones	Black	1.8	0	95.4	0
DeShawn Glover	Black	4	1	95	0
Tyrone Thomas	Black	3.9	0.6	94.7	0
Terrell Turner	Black	4.4	0	94.4	0

Darnell Jackson	Black	5.7	0	94.3	0
Terrell Watkins	Black	5	0.8	93.1	0.4
Trevon Williams	Black	7.1	0	92.9	0
Darius Haynes	Black	6	0.7	92.7	0
DeAndre Wilkins	Black	5.3	0.3	92.3	0
Darnell Haynes	Black	7.5	1.1	91.4	0
DeShawn Ware	Black	5.4	0	91.2	0
DeAndre Scott	Black	5.8	0.4	91.2	0
Trevon Johnson	Black	0.9	0	90.9	0
Tyrone Jones	Black	9.2	0	90.8	0
Jalen Washington	Black	6.9	0	90.8	0
Darius Davis	Black	9.3	0	90.7	0
Darnell Alexander	Black	8.3	0.5	90.4	0
DeShawn Anthony	Black	3.5	0	90	0
Demetrius Jackson	Black	10	0	90	0
Darnell Davis	Black	11.8	0	88.2	0
Terrell Davis	Black	10.9	0	88.2	0.9
Jamal Coleman	Black	7.5	0.5	88	4
Tyrone Johnson	Black	8.5	0	87.7	0
Darius Washington	Black	11.8	0.6	87.6	0
Marquis Harris	Black	6.5	5	87	0
Malik Johnson	Black	5.5	0	86.4	6.4
Maurice Brown	Black	13.8	0	86.2	0
Tyrone Harris	Black	11.5	0.3	85.5	0
DeShawn Johnson	Black	13.6	0	85	0
DeAndre Davis	Black	12.7	1	85	0
Terrell Ware	Black	6	1.8	84.5	1.8
Andre Harris	Black	13.1	1.5	84.2	0
Jamal Williams	Black	10.5	1.1	84.2	1.1
Darnell Mitchell	Black	15.4	0	83.9	0
Darnell Carter	Black	10.3	0	83.8	0
Terrance Terrell	Black	13.5	1.2	83.5	0
Terrell Scott	Black	12.5	0.2	83	0
Terrance Johnson	Black	17.5	0	80.8	0
Andre Johnson	Black	19.3	0.2	80.4	0
Terrell Washington	Black	12.3	0	80.3	0
Demetrius Johnson	Black	14.5	0.5	79.1	0
Darryl Willis	Black	20	0	79	0
Dominique Richardson	Black	18.4	2.7	78.9	0
Darius Miles	Black	20.5	0.5	78.6	0
Darius Willis	Black	13	0	78.3	0
Dominique Brown	Black	16.2	0	77.2	0
Darius Bryant	Black	20	1.1	77.2	0
Trevon Grant	Black	20	1.7	77.1	0
Trevon Henry	Black	20.6	2.1	76.8	0
Reginald Brown	Black	13	8.5	76.5	0
Marquis Williams	Black	15	0.8	75.7	0
Dominique Walker	Black	21.8	1.6	75.5	0
Malik Hawkins	Black	15.9	0.3	75.3	8.3
Tyrone Dorsey	Black	25	0	75	0
Terrance Robinson	Black	16	0.2	73.8	0
Darius Byrd	Black	20.4	0	73.5	0
Malik Williams	Black	0.3	0.8	73.3	19.7
Jalen Walker	Black	27.1	0	72.3	0

Trevon Scott	Black	25.8	0	71.7	0
Maurice Miles	Black	25.2	0.5	71.5	0
Malik Mitchell	Black	6.7	0	71	14
Jamal Johnson	Black	6	0	71	3
Xavier Brown	Black	16.2	6.9	70.3	0
Dominique Jones	Black	22.7	4.5	70	0
DeAndre Mathis	Black	16.3	3.7	69.7	0
Maurice Davis	Black	29	0.6	69.4	0
Terrell Thomas	Black	8.3	8.3	69.2	8.3
Reginald Coleman	Black	33.3	0	66.7	0
Jalen Neal	Black	20	0	65.8	0
Jalen Harris	Black	17.8	2.8	65	0
Maurice Thomas	Black	27	1.3	64.3	0
Darryl Brooks	Black	28.9	7.1	62.1	0
Reginald Davis	Black	39.2	0	60.8	0
Malik Robinson	Black	14.4	0	60.6	18.9
Marquis Mitchell	Black	17.7	3.1	60.4	0
Terrance Woods	Black	39.3	0	60.4	0
Jalen Johnson	Black	10	0	60	3.3
Demetrius Fields	Black	23.5	2.4	60	0
Dominique Simmons	Black	27.7	11.2	59.6	0
Jalen Thomas	Black	26.8	4.5	59.5	0
Darryl Watkins	Black	39.1	0	57.7	0
Jalen Carter	Black	36	0	57.5	0
Xavier Scott	Black	37.8	0.6	56.7	3.3
Xavier Willis	Black	20.7	20	56.4	0
Willie Davis	Black	40	1	56	0
Malik Neal	Black	16.3	0	55.8	14.2
Xavier Brooks	Black	28.1	0.8	55	0
Dominique Alexander	Black	30.6	12.1	55	0
Willie Brown	Black	37.8	0.4	54.8	0.9
Darryl Williams	Black	28	0	54.5	0
Willie Jones	Black	39	2.5	54.5	0
Willie Williams	Black	43.3	0	54.3	0
Dominique Matthews	Black	34.7	8.8	53.5	0
Andre Miles	Black	35.8	9.2	52.3	0
Xavier Davis	Black	44	0.3	49	0
Darryl Brown	Black	44.4	0.6	47.8	0
Darryl Davis	Black	53.2	0	45	0
Willie Singleton	Black	46.2	0	43.8	0
Reginald Turner	Black	45	5.6	40.8	0
Jalen Holmes	Black	33.6	0	40.5	0
Darryl Walker	Black	57.3	0.7	40	0
Willie Nixon	Black	71.4	0	13.6	0
Basir Albaf	Arab	0	0	0	99.2
Botros Ahmed	Arab	0	0	0	98.4
Sami El-Amin	Arab	0	0	1.7	97.8
Salah Darzi	Arab	0	0	2.2	97.8
Abd El-Mofty	Arab	0	0.5	0.9	97.7
Sharif Abdullah	Arab	0	0	2.9	97.1
Shahnaz Hussain	Arab	0	0	0	96.8
Duha El-Amin	Arab	0	0	1.5	95.8
Shams El-Amin	Arab	0.1	0.1	3.3	95.6
Ibrahim El-Hashem	Arab	0	0	1.8	95.5

Mahdi Albaf	Arab	0	0	1.8	94.7
Bakr Abdullah	Arab	0	0	0	94.5
Husain Sultan	Arab	0	0	0	94.4
Sajjad Ahmed	Arab	0.6	0	1.2	94.1
Fayiz Muhammad	Arab	0	0	1	94
Ghassan Ahmed	Arab	6.2	0	0	93.8
Ghayth Abdullah	Arab	0	0	4.7	93.6
Ramadan Muhammad	Arab	0	0	4.4	93.3
Maalik El-Ghazzawy	Arab	0	0	1.9	93.1
Hafeez Saab	Arab	0	0	3	93
Tarik El-Amin	Arab	0	0	5	93
Abbas Abdullah	Arab	0	0	4.2	92.9
Imad Zaman	Arab	0	0	1.4	92.9
Mohammed Ahmed	Arab	0	0	3.8	92.5
Jabr Hussain	Arab	5.9	0	1.8	92.4
Hikmat Ahmad	Arab	1.2	0	0	92.2
Bahadur Abdullah	Arab	0.7	0	0	92.1
Al-Amir Bousaid	Arab	0	0	0.3	92.1
Shadi Bousaid	Arab	0	0	0	91.7
Jalal El-Amin	Arab	0	0	1.9	91.5
Nasim Abdullah	Arab	0	0	2.6	90.9
Salil Albaf	Arab	2.1	0	0.7	90.7
Hakim Ajam	Arab	0	0	8.7	90.7
Boulos Amjad	Arab	1.2	3.8	1.9	90.6
Baqir Ali	Arab	3.3	0	0.8	89.2
Mohammed Boulos	Arab	0	0	11.2	88.8
Bahij Nejem	Arab	0	0	0.9	88.6
Zahi El-Mofty	Arab	0	0	0.7	88.6
Gafar Hakim	Arab	0	0	2.9	88.6
Hussein Darzi	Arab	0.6	1.8	3.2	88.2
Basir Muhammad	Arab	0	2.1	8.6	88.2
Sa'Di Albaf	Arab	0	6.7	3.7	88
Mukhtar Amjad	Arab	0.5	0	6.5	87.8
Tahir El-Amin	Arab	0	4.6	2.4	87.6
Yuhanna El-Amin	Arab	0	0	6.2	86.9
Aamir Abujamal	Arab	0	0	0.8	86.7
Husain El-Mofty	Arab	10.9	0	0.9	86.4
Fadl Nejem	Arab	0	0	0	85.7
Halim Zaman	Arab	0	0	2	85.5
Imran Hakim	Arab	7.7	1.5	1.5	85.4
Samir Abdulrashid	Arab	0	0	1.1	84.6
Ihsan El-Mofty	Arab	0	0	0	84.5
Tarek Saqqaf	Arab	0.7	0	6	84
Abdul-Aziz El-Mofty	Arab	0	0	1.6	83.2
Wadud Hakim	Arab	1.2	0	13.8	82.5
Shukri Saqqaf	Arab	0	0	3.8	82.3
Yaser Karimi	Arab	0	0	3.2	81.6
Fakhri Ali	Arab	0.1	0	5.3	80.8
Nabil Saab	Arab	0.6	0	7.8	80.6
Ziauddin Muhammad	Arab	0	0	1.2	80
Rayyan Albaf	Arab	0	0	5	79.3
Rasul Ajam	Arab	0	0.3	1.5	78.8
Nour El-Ghazzawy	Arab	1.5	0	3.1	78.5
Rifat Alfarsi	Arab	0	0	6.7	78.3

Sajjad El-Amin	Arab	0	0	5	78.3
Sa'Di El-Ghazzawy	Arab	0.7	0	8	77.3
Fayiz Samara	Arab	1.5	0	2.3	76.2
Aali Hussain	Arab	0	11.1	1.1	75
Imran Mohammed	Arab	1.1	0	6.7	74.4
Nizar Kader	Arab	0	0	2.8	73.9
Jaffer Bousaid	Arab	6.9	0	1.2	73.8
Jafar Sultan	Arab	0.3	0	17.6	73.2
Shafiq Samara	Arab	0.9	0	16.8	73.2
Fayiz Nejem	Arab	0	0.3	2.6	72.4
Salim Kader	Arab	0	0	10.4	72.1
Wafi Sultan	Arab	0	0	3.7	71.6
Husni Zaman	Arab	0	0	18	71.3
Adam Ahmad	Arab	7.4	5.2	7.4	71.0
Khaled Samara	Arab	0	3.3	14.7	70
Rasheed Zaman	Arab	2.7	0.7	22.7	70
Fakhri El-Mofty	Arab	1.8	0.3	12.9	68.8
Sameer Sultan	Arab	6.2	0	9.6	68.5
Guda El-Mofty	Arab	0	11	7.5	66.5
'Abbas Nagi	Arab	0	0	15.5	65
Adnan El-Mofty	Arab	0	0	8.3	64.2
Zaki Karim	Arab	1.1	0	20.3	63.9
Mis'Id El-Ghazzawy	Arab	0	0	0	63.3
Nurullah Nejem	Arab	0	1.1	10.8	61.9
Latif El-Mofty	Arab	0.5	3.2	29.2	61.6
Safi Boulos	Arab	0.4	7.7	0.4	61.5
Tayeb Kader	Arab	3.8	0	21.8	59.8
Waheed Bousaid	Arab	1.5	0	14.4	58.5
Mansoor Amirmoez	Arab	0	21.2	5.6	58.1
Dawud Karim	Arab	0	1.2	35.6	52.9
Tal'At Tawfeek	Arab	7.1	0	20	46.4
Murtaza Nagi	Arab	0.4	0.7	4.6	42.5
Ayman Amirmoez	Arab	0	28.1	0	41.9
Rusul Samara	Arab	1.8	5.9	14.5	41.4
Rais Nagi	Arab	0	0.1	1.9	40
Wafi Kader	Arab	2.5	0	23.8	33.8

References

Abdulrahman M. Al-Sayed, Diane S. Lauderdale, and Sandro Galea. Validation of an arab names algorithm in the determination of arab ancestry for use in health research. *Ethnicity & Health*, 15(6):639–647, 2010.

Marianne Bertrand and Sendhil Mullainathan. Are emily and greg more employable than lakisha and jamal? a field experiment on labor market discrimination. *American Economic Review*, 94(4):991–1013, 2004.

Daniel M Butler and David E Broockman. Do politicians racially discriminate against constituents? a field experiment on state legislators. *American Journal of Political Science*, 55(3):463–477, 2011.

Roland Fryer and S. Levitt. The causes and consequences of distinctively black names. *Quarterly Journal of Economics*, 119(3):767–805, 2004.

Jens Hainmueller, Jonathan Mummolo, and Yiqing Xu. How much should we trust estimates from multiplicative interaction models? simple tools to improve empirical practice. *Political Analysis*, In Press, 2018.

Winston Lin and Donald P Green. Standard operating procedures: A safety net for pre-analysis plans. *Science*, 343(6166):30–1, 2015.

Ariel R. White, Noah L. Nathan, and Julie K. Faller. What do i need to vote? bureaucratic discretion and discrimination by local election officials. *American Political Science Review*, 109(1):129–142, 2015.

David L Word, Charles D Coleman, Robert Nunziata, and Robert Kominski. Demographic aspects of surnames from census 2000. *Technical Report for the U.S. Census Bureau*, 2008.